1. FUSION ENERGY RESEARCH

Since the late eighties, the Fusion Energy Research Program here at Berkeley Lab had been focused on the construction of the Induction Linac Systems Experiments (ILSE), which would have addressed many important accelerator issues at a scale significantly close to that of a fusion-energy driver. But nationally, funding for fusion has been decreasing for more than a decade. Last year the decrease was precipitous (about 30% between FY95 and FY96).

In large measure, the situation in inertial fusion research at Berkeley Lab mirrors the more global financial picture in the Department of Energy's Office of Energy Research. It is possible to exploit existing facilities to do good science, but it is extremely difficult to get new projects, particularly large new projects, approved. The Superconducting Super Collider was started and then canceled. Two proposed magnetic-fusion devices, the Burning Plasma Experiment (approximately $2B) and the Tokamak Physics Experiment ($0.5–1B) were both canceled after a large amount of design effort and technology development. The future of the International Thermonuclear Experimental Reactor is uncertain.

Under these circumstances it was not possible to build ILSE rapidly enough to make programmatic sense without eviscerating our scientific work in order to support it. Moreover, advances in science and technology, particularly advances in our own program, led us to believe that we could build a substantially better machine than ILSE—a machine that would, at acceptable cost, enable some target physics as well as accelerator physics.

Because of this situation, we restructured our research to emphasize small-scale experiments, theory and computation, and enabling technologies. The goal of our program is to provide the science and technology base needed to build accelerators that deliver far more energy per dollar than previous designs. This has been a long-standing interest of ours, because we have known all along that powerplant economics will be as important as technical feasibility in the ultimate success of fusion energy. The basic structure of the development path leading to a heavy ion fusion power plant was always attractive economically. We have estimated that this path to fusion energy would cost 25 to 50% as much as the path to magnetic fusion energy. Nevertheless, accelerator cost reductions are desirable, since full-scale fusion power plants based on today's concepts are too expensive and, in today's funding climate, the development path is also too expensive. Although heavy-ion fusion looks better than the mainstream magnetic approaches (tokamaks and stellarators) in this regard, all kinds must do better in terms of economics as well as science.

In June and July 1996, a subpanel of the Fusion Energy Sciences Advisory Committee reviewed and endorsed the new direction of our program. The

Subpanel specifically stated that a central goal of the Program should be to enable the construction of a multi-kilojoule accelerator facility, beginning near the turn of the century, to establish the science and technology needed to build a full-scale heavy ion fusion driver.

As the next step after the multi-kilojoule facility, and undoubtedly not at the Berkeley site, we ultimately plan to build a full scale fusion accelerator that can be used for target physics, fusion chamber R&D, and finally as the driver for a demonstration power plant.

Long-Range Planning for a Successful Driver, and How This Guides Our Program

In most studies of inertial fusion power plants, the driver is the single most expensive element. It is usually expected to cost of the order of $1B, which is roughly half the total cost of the power plant. A driver cost of $500M gives very good projected economics, and a driver cost of a few hundred million dollars would be phenomenal. Driver efficiency is also important. An efficiency of 25% gives acceptable economic projections. Considering this situation, we have asked ourselves two important questions:

1. Can we build a full-scale driver for approximately $250M?
2. Can we build a driver with an efficiency approaching 50%?

Today, the answer to both questions is “no,” but existing accelerators provide some guidance toward changing the answer to “yes.”

Most major accelerators have been built for a cost of $100–$400k per meter. Some induction machines have been even more expensive. Many of our accelerator designs produce ions with a kinetic energy of 10 GeV, and we usually assume that we can achieve an accelerating gradient of approximately 1 MV/m. These numbers lead to an accelerator length of the order of 10 km and a cost exceeding $1B. We must reduce the length or the cost per unit length—preferably both. There are a number of ways to do this:

1. Reduce the accelerator voltage. To do this, one can either increase the charge of the ions or reduce their kinetic energy. Reduced kinetic energy, for fixed ion range, requires a reduction in ion mass so both methods increase the charge-to-mass ratio. If one pushes this strategy very far, space charge neutralization is needed at the final focus and preferably in final longitudinal compression also.
2. Decrease the transverse size of the accelerator.
3. Increase the acceleration gradient.
4. Decrease the length by recirculating.
5. Improve technology and fabrication methods. The intrinsic cost of materials is usually low; turning them into finished parts and then assemblies is what drives up the cost.

6. Reduce the target's energy requirements. Today most target designs require 5 to 10 MJ of energy to achieve adequate energy gain. Innovations such as fast ignition or direct drive will, if successful, reduce the requirements by more than a factor of two.

We have made substantial theoretical and experimental progress in most of these areas during the last year, as you will see in this chapter and hear in the review talks. In this chapter we include four technical sections that illustrate some of our work in the first two areas listed above.

The first section on ion source research emphasizes high current density. High current density leads to small transverse dimensions.

The second section describes experimental work on beam combining. The number of beams required to minimize accelerator size and cost varies with kinetic energy along the accelerator. Combining allows the number of beams to vary so that one can choose the optimum number at different points in the accelerator. Unfortunately, combining also increases emittance. Our experiments and their detailed analyses will enable us to make the proper design choices between these competing effects. (We have also initiated a theoretical effort on beam splitting, another technique that allows the number of beams to be optimized for various points in the same machine.)

The third section describes work on plasma channel focusing. As noted above, neutralization is essential to large reductions in accelerator voltage and length.

The fourth section of the report describes solenoidal transport. Solenoidal transport scales favorably at low kinetic energy, so this work is related to the experiments described in section 3. The improvements in current density from the ion source described in section 1 are also essential if we are to optimize the solenoidal transport option.

Reduction of target requirements, the sixth method of cost reduction listed above, plays an important role in our program planning. For indirectly driven targets, much of the physics is independent of the type of driver. Existing lasers, and to a smaller extent existing accelerators, have already put to rest many of the important issues related to indirectly driven targets. The National Ignition Facility will essentially complete the needed data base for indirectly driven targets for both lasers and accelerators. It will also complete the needed data base for directly driven laser targets. One major area of target design remains nearly unexplored experimentally, namely direct drive for ions. This situation is illustrated in Figure 1-1. Experimental exploration of this area of target physics will be an important goal of the aforementioned new accelerator facility.

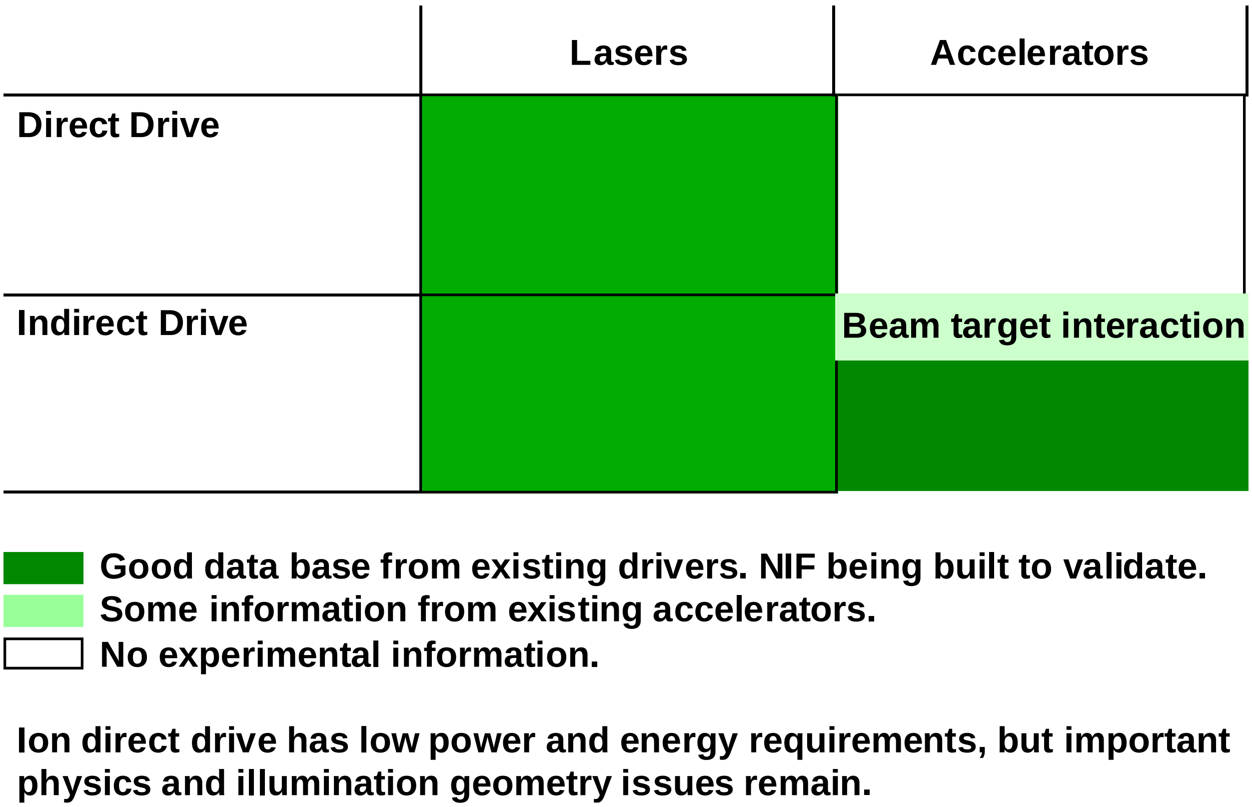

Figure 1-1. *Direct drive targets for accelerator use represent the major completely unexplored area of target design. Exploring this science is a major goal of the new multi-kJ accelerator that we would like to build in the near future.*

In this year's chapter, we do not present a separate section on theory and numerical modeling. Nevertheless they are some of our most important activities. We have begun an effort to follow a computational beam ensemble from its birth at the anode surface to its absorption by a target. This endeavor will provide critical information needed for decisions concerning the design of any large-scale accelerator for heavy-ion fusion research. Together with collaborators at Lawrence Livermore National Laboratory and Princeton Plasma Physics Laboratory, we are adapting and modernizing a number of particle simulation codes (and other types of codes) to study areas such as multiple-beam injectors, long distance transport in electric and magnetic focusing lattices, final focus lattice design, and transport in the target chamber. Results from these codes are needed for design decisions such as error tolerances and resultant emittance growth within the accelerator, optimal numbers of multiple beamlets, the efficacy of beam combining and/or splitting, and tradeoffs between different types of focusing schemes. This effort is related to Berkeley Lab's acquisition of the National Energy Research Scientific Computing Center (NERSC) and our subsequent efforts to port codes to their massively parallel Cray T3E supercomputer.

Programmatic Issues

We had hoped to complete a conceptual design of the new accelerator by the end of January 1998. Meeting this date would enable construction to begin in FY2000. We are now facing two challenges which make it difficult to meet that goal. The first challenge is staffing. We are deficient in conventional accelerator physics and technology. The problem is even more severe in engineering. We do not have the money to hire the needed engineers even if they were available. Without them we cannot obtain the needed engineering data and produce a credible conceptual design. In summary, we do not have the proper balance on our team, nor do we have the flexibility to achieve the balance unless some funding relief is in the offing.

The second challenge is partly technical; it involves the choice of technology for the next accelerator and for a driver. There are two principal classes of options: linacs and recirculators. (The linacs can be further divided into two classes that are somewhat different. The first class uses a number of small beams to bring the space charge problems into an acceptable range. The second class uses a smaller number of large beams and relies more heavily on neutralization and advances in accelerator technology.) We have developed three preliminary designs for the new accelerator. The designs are currently rather superficial, but they do elucidate the issues. The designs are a multi-beam linear configuration using quadrupole focusing; a linear configuration using a single solenoidally focused beam; and a recirculating system.

The decision between the linac options and the recirculator is also at least partly institutional. Livermore and Berkeley collaborate on both the linacs and the recirculator, but the recirculator is more strongly identified with Livermore. Neither the people who work on the recirculator nor the people who work an the linacs would like to see the selection process eliminate their option. We hope to make the decision by October of this year, although the date may slip to the end of the calendar year. To get the support of both laboratories, the decision must be made on firm technical grounds. We have made good progress toward the decision, but several critical issues must still be resolved. The October date appears optimistic, but we are confident that we will be able to make this decision on firm technical grounds within a year.

Roger Bangerter, Program Head

Ion Source R&D for Fusion Energy

Ion source performance—current density in particular—has an important direct impact on the cost-effectiveness of heavy-ion inertial fusion energy system for two reasons. First, high current density leads to small transverse dimensions and therefore a cheaper accelerator. Second, it is essential for optimization of solenoidal transport, an attractive option for the low-kinetic-energy portion of a driver as described later in this chapter.

Performance and Progress

The ion source that we had developed in previous years for the Elise injector has a 6.7" diameter aluminosilicate emitter producing 0.8 A of K^+ ions at a current density of 3.5 mA/cm^2. The emittance is < 1.0 mm-mrad. The corresponding line charge density is approximately 0.25μC/m, similar to that required in a fusion driver. Although this design met the beam current and emittance requirements for the since-abandoned Elise project, the cost of such a single-beam injector (approximately $2M per copy) does not scale well for a fusion driver that may require 100 beams or so. One way to minimize the overall injector cost is to develop compact multiple-beam injectors using high-current-density ion sources.

We have made significant progress in improving the quality of the aluminosilicate coating on the tungsten substrate. Poor coating uniformity can adversely affect the beam emittance. By using smaller particles, as well as suction to "prime" the porous tungsten substrate, we found that the aluminosilicate-silicate can be spread evenly on the tungsten surface without exposing uncoated areas.

We have built a triode extraction system to test 2-cm diameter aluminosilicate-silicate emitters. The extraction system was designed using the WOLF code and independently verified by the EGUN code. So far, experimental results (see Figure 1-2) have reached a current density of 14 mA/cm^2 of K^+ ions, limited by the extraction voltage according to space-charge-limited flow scaling. As shown in Figure 1-3, the current density is already quite uniform and is expected to further improve with better quality surface coating. As of this writing, we have just completed upgrading the high voltage pulser to deliver higher extraction voltage. Thus further increase in the current density is expected in the near future.

At 14 mA/cm^2 the current density is already high enough to allow the construction of a compact (approximately 1 m^2 of emitters) injector for a full-scale fusion driver.

Recent progress in producing chopped ion beam bunches for the National Spallation Neutron Source progress (by K.-N. Leung in AFRD's Ion Beam Technology program; see Chapter 6) has encouraged us to re-examine the feasibility of using plasma sources for heavy ion fusion. The key here is to produce fast-rise-time (< 1 μs) beam pulses suitable for induction linacs. An experiment is being set up with the Ion Beam Technology Program.

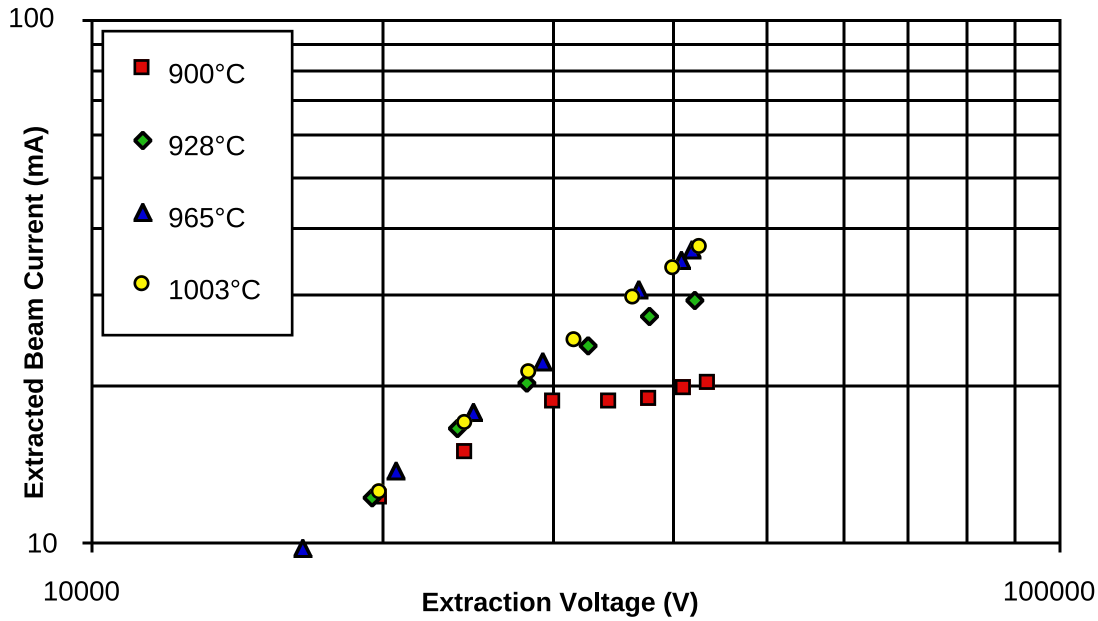

Figure 1-2. *Extracted beam current as a function of extraction voltage for the 2-cm K^+ ion source.*

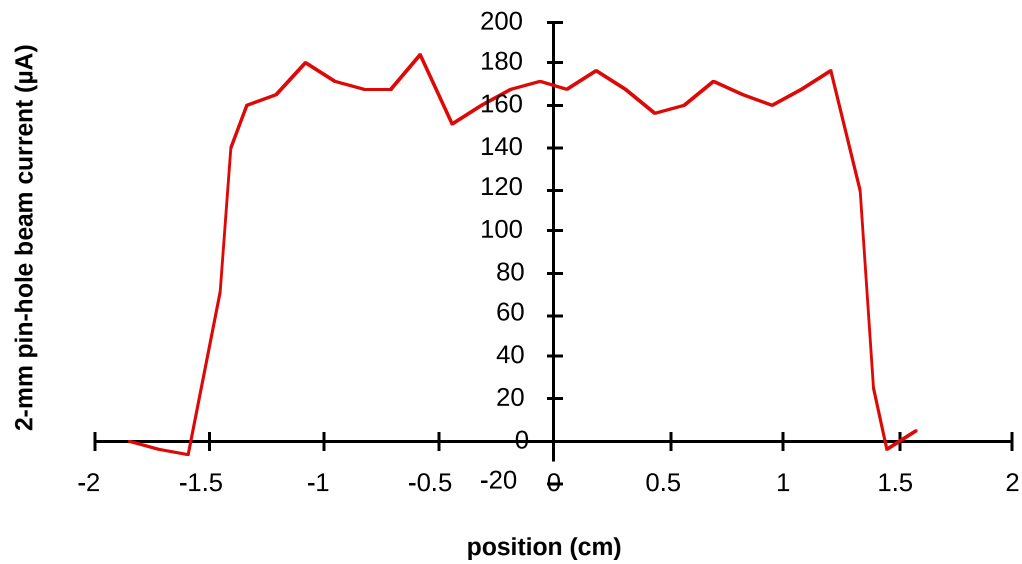

Figure 1-3*. K^+ beam profile from the 2-cm-diameter source.*

Reported by Joe Kwan

Beam Combining Experiment

Transverse beam combining is a cost-saving option employed in many designs for induction-linac heavy-ion-fusion drivers. The resultant transverse emittance increase must be kept minimal so that the beam remains focusable at the target. A prototype combining experiment has been built, incorporating the MBE-4 apparatus.

Experiment Design

Four new sources each produce up to 5 mA Cs^+ beams at 160 keV. The ion sources are angled toward each other, so that the beams converge. Focusing upstream of the merge point consists of four quadrupoles and a final combined-function element (quadrupole and dipole). All lattice elements are electrostatic. Due to the small distance between beams at the last element (~3-4 mm), the electrodes are a cage of tungsten rods, approximately parallel to each of the four beam trajectories. There are 71, one mm diameter wires, each at different voltage set to produce a combined quadrupole and dipole. The beams emerge into the 30-period transport lattice of MBE-4 where the subsequent evolution of the phase space and emittance growth from merging can be measured.

The emittance growth from beam merging arises from emittance dilution associated with the efficacy of packing the beams in phase space (single particle dynamics) and also from the conversion of the potential energy in the space charge field of the merged beams to kinetic energy. In a 2D approximation, these two growth mechanisms add in quadrature. The space charge component is

$$\Delta \varepsilon_n^2 = \frac{2\lambda q r^2 \delta f}{\gamma m c^2}$$

where δf is a dimensionless geometry dependent factor, r is the merged beam radius, and λ is the line-charge density. Even though λ in this scaled experiment is much smaller than in a driver, the relative contribution to emittance growth from space charge is — by design — commensurate with what might occur in a full scale driver combiner.

Figure 1-4 shows a CAD view of the lattice elements of the combiner, and the photo insets below show some of the combiner hardware. All of the apparatus is installed in the region formerly occupied by the MBE-4 matching section.

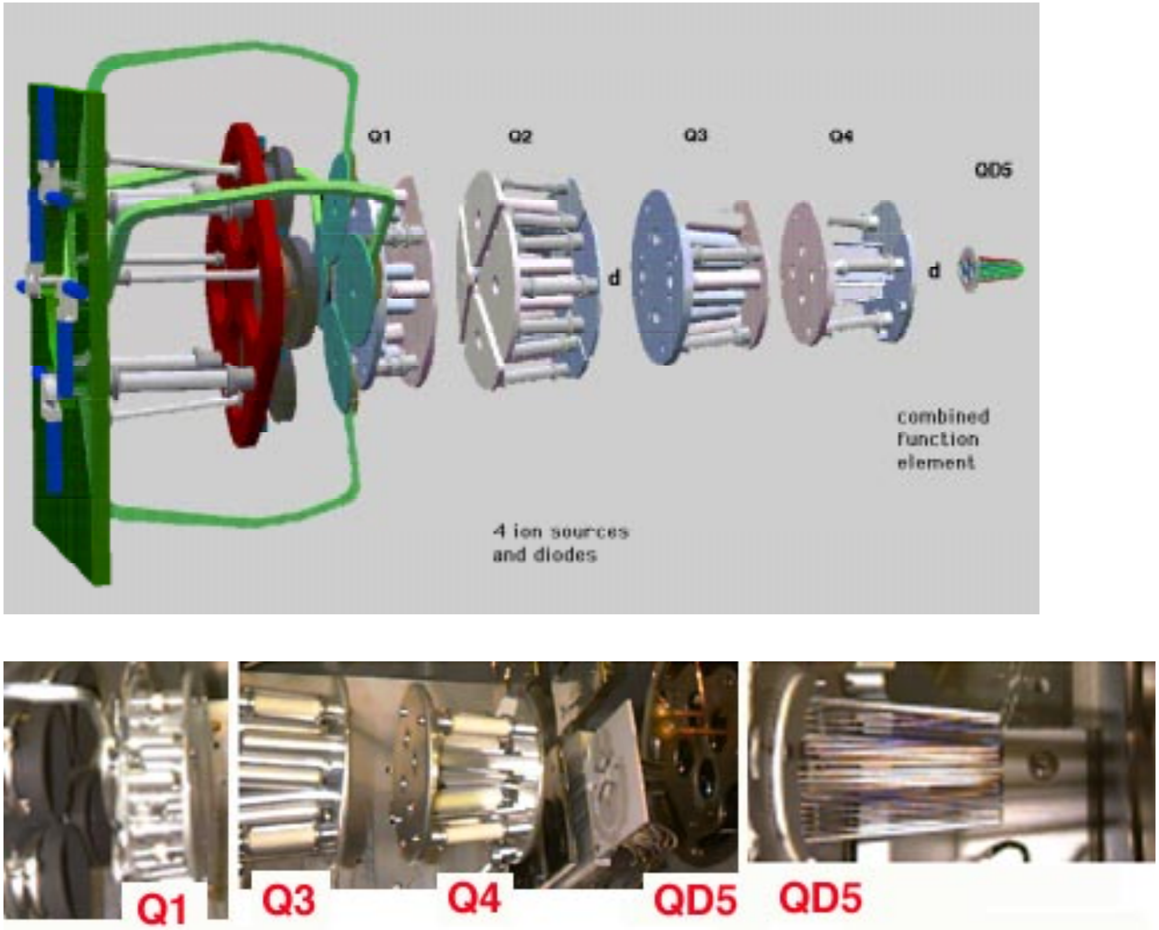

***Figure 1-4**. A CAD view of the experiment with photos of some of the beam focusing elements. The array manipulates Cs^+ at about 160 keV and 3–5 A per beam. From left, four ion sources and associated diodes (black disks like hockey pucks in the photo, color-coded red in the diagram) send beam into the successive quadrupole arrays, where they are converged at an angle of about 100 milliradians. QD5, at far right, is the combined-function element. The first two phase-space diagnostic stations are denoted by "d" in this figure.*

2D particle-in-cell (PIC) computer simulations of the beams traveling through the combiner apparatus and the MBE-4 transport line (see Figure 1-5) have predicted the emittance growth due to merging as well as the beam loss.

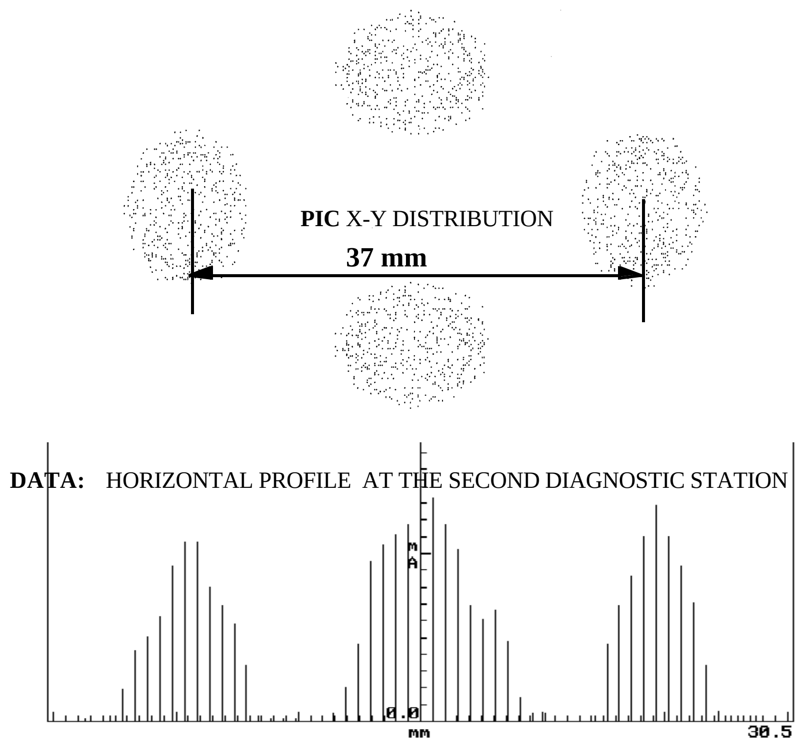

Figure 1-5. *The top scatter plot is a 2D particle-in-cell (PIC) simulation of the x-y distribution of the four beams after traversing the first four quadrupoles, or, just upstream of the combined function element (QD5 in Figure 1-4). The histogram below it shows preliminary experimental data: a horizontal profile of the beams measured with a vertically oriented slit and Faraday cup diagnostic at the same location as the PIC simulation. The beam positions are in fair agreement, and the asymmetries in the profile shapes are due to various misalignments whose corrections are now being undertaken using the remotely articulable sources and second quadrupoles (Q2).*

3D field calculations and multipole decomposition of the wire cage focusing and bending element were used to study the unusual design of QD5. This was followed by calculations of the emittance growth due to the non-linear field components, which was relatively small compared to the final merged beam at the end of the accelerator.

The main body of experimenting and interpretation is underway. This will provide useful insight into the physics and technical feasibility of merging beams. Preliminary results were reported at the IAEA Fusion Energy conference and the APS Plasma physics conference in October and November 1996.

During the past year, the experiment was refined considerably, and beam manipulations began. The source steering and wire cage HV divider circuits were modified, and changes were made to eliminate the influence of secondary electron space charge, which introduced an unwanted time dependence and halo. The phase spaces and profiles of the beams have been measured at two diagnostic stations upstream and one downstream of the wire cage. Above, we saw a beam profile at the diagnostic station just upstream of the wire cage, measured in the horizontal plane, compared to a PIC simulation of the x-y spatial distribution of the beam at that point.

Most of the ions are transmitted through the cage, which functions (with no evidence of voltage breakdown) as a combined dipole and quadrupole, just as designed. Upstream, the four beams show various misalignments whose corrections are now being undertaken using the remotely articulable sources and second quadrupoles. The next work will include measurements of the beam's evolution at several more downstream diagnostic stations and comparisons with PIC simulations.

Reported by Peter A. Seidl

Plasma Channel Focusing of Heavy Ion Beams

Final focusing of ion beams and their propagation in a target chamber are crucial for heavy-ion-beam-driven fusion power production. The method and technical realization of these operations have strong impacts both on the design of the accelerator and the layout of the target chamber. Most studies today are based on ballistic focusing by magnetic quadrupole lenses. With ballistic focusing, the focal spot size depends on the current and shape of the beam pulse that is delivered by the driving accelerator. To reduce space charge blow-up of the driver beam in the final focus region the total beam current is divided into several beams, each with an individual final focusing system and a separate beam port in the target chamber.

There are several possible alternative methods of focusing and transporting the beam. This report describes one concept that offers a number of desirable features: plasma channel transport.

Plasma Channel Transport

Plasma channel transport is illustrated in Figure 1-6. The beam is injected into a current-carrying plasma channel. Collisions cause the beam ions to be stripped to a high charge state, and the plasma neutralizes the space charge and current. Plasma lens focusing significantly relaxes the requirements on emittance and energy spread of the beam. Also, the magnetic field of the channel current confines the beam ions within the channel. Channel transport has the added advantage of physical simplicity because the beam can be treated as an ensemble of independent particles in the external focusing field. A further advantage of this concept is insensitivity to beam current and pulse shape.

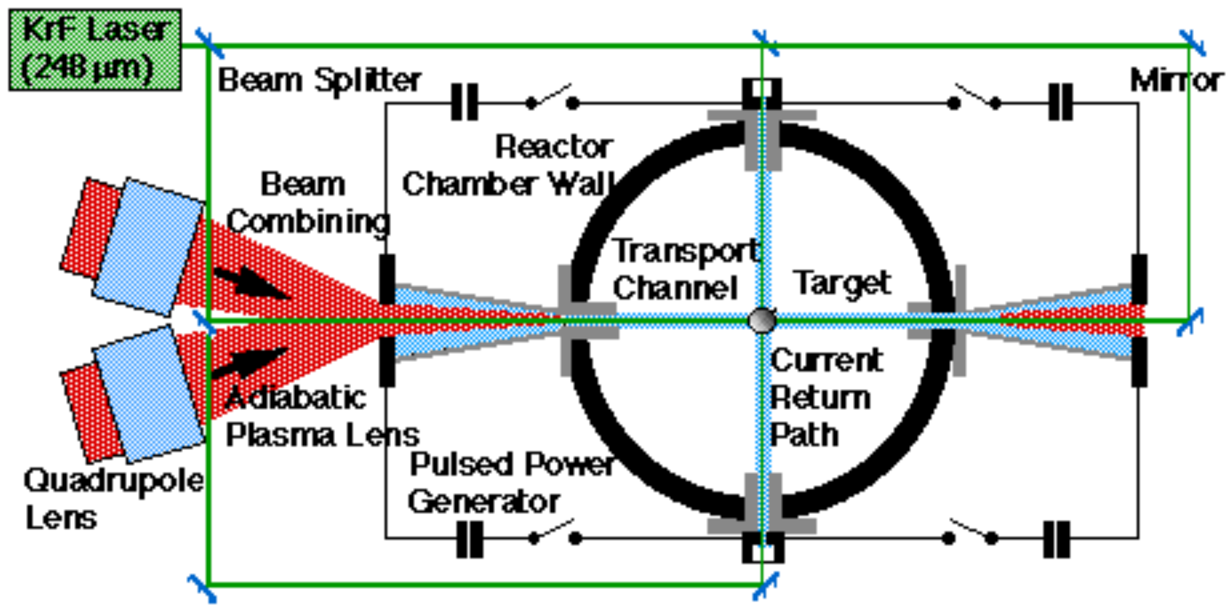

Figure 1-6. *Channel focusing of heavy ion beams to an inertial fusion target.*

A pulsed laser is used to "designate" the current channels to and from the fusion target. A high voltage applied along the channel breaks down the gas and drives several tens of thousands of amperes through the channel. The theta-directed fields generated by the channel current are sufficient to confine the heavy ion beams and precisely guide them to the fusion target. Largely because light ion fusion requires high beam currents, this approach was originally considered by researchers studying the light ion approach to inertial fusion.

We are performing two experiments at LBNL to investigate this concept for heavy-ion fusion. One experiment studies the physics of the adiabatic focusing of an intense potassium beam generated by the 2-MeV Injector. The other experiment studies the formation and interactions of free-standing Z-discharge channels at parameters that are appropriate for driver-quality heavy ion beams.

Adiabatic Focusing of an Intense Ion Beam

Plasma lenses for heavy ions have been developed and investigated for several years at the German heavy ion laboratory GSI to concentrate high beam intensities onto small focal spots. The most promising results were achieved with a wall-stabilized discharge. [1] The typical beam energy for the GSI experiments was 11.4 MeV/amu. Discharge current and length of the lens were matched to produce less than half a betatron oscillation of the ions in the lens; that is, it was a thin lens. The focus was formed a short distance downstream of the lens and detected by a combination of a plastic scintillator and a streak camera. A wall-stabilized discharge has demonstrated better linearity of the focusing fields than a z-pinch discharge and provided the

[1] E. Boggasch, A. Tauschwitz, H. Wahl, K.-G. Dietrich, D.H.H. Hoffmann, W. Laux, M. Stetter, R. Tkotz, "Plasma lens fine focusing of heavy-ion beams", Appl. Phys. Lett. 60 (1992) 2475-2477.

additional opportunity to taper the discharge tube and increase the focusing strength along the lens. Tests of a tapered discharge showed that a beam with 10 mm initial diameter can be focused to a spot of 0.16 mm. [2]

For inertial confinement fusion final focus, the plasma lens must work as a thick lens; that is, the particles perform one or more complete betatron oscillations inside the discharge. The particle trajectories in the lens are not necessarily coherent. The broader the momentum and charge state distribution in the beam, the more the trajectories deviate from the coherent case, and the more oscillations the particles perform in the lens. In the extreme case the beam fills the aperture homogeneously everywhere within the lens. Focusing can only be accomplished by increasing the focusing strength along the lens which can be easily achieved by tapering the discharge tube diameter. To avoid beam losses the tapering must be done adiabatically, which means that the change in betatron wavelength per betatron period must be small.

An experiment was designed to investigate adiabatic focusing experimentally using the 2 MeV beam from the LBNL-ESQ injector. The schematic layout for this experiment is sketched in Figure 1-7. The beam traverses a two stage differential pumping system that separates the discharge gas of approximately 1 Torr from the accelerator vacuum of 10^{-6} Torr. The discharge tube radius is tapered from 10 mm to 2.5 mm to reduce the beam radius from 5 mm to 2.5 mm. The intensity distribution in the focused beam is scanned by a pinhole close to the end of the discharge and measured with a Faraday cup.

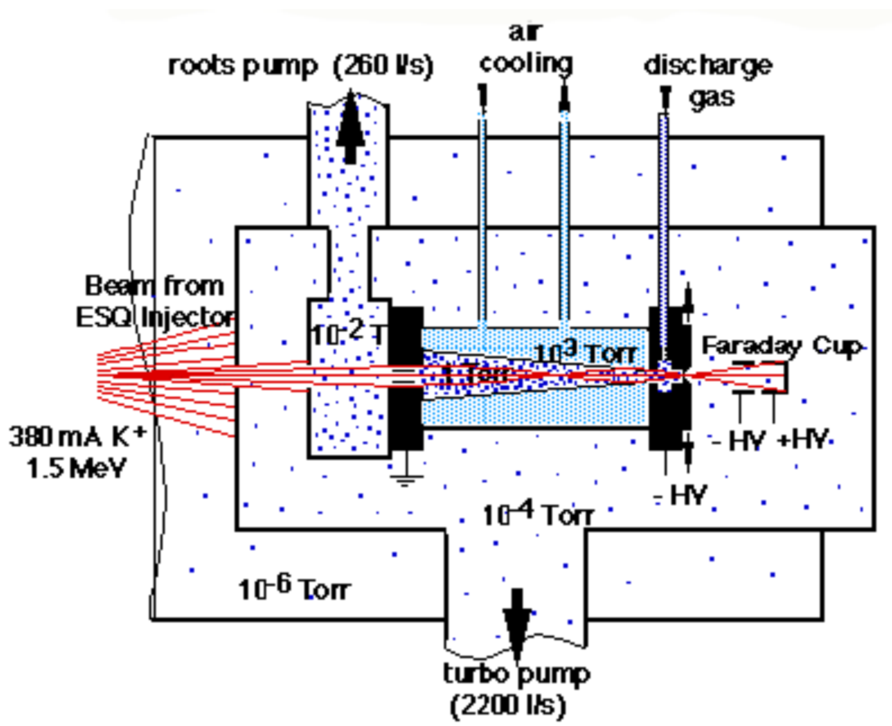

Figure 1-7. *Schematic of the adiabatic focus apparatus mounted after the matching section of the 2 MV Injector. The ion beam enters from the left and passes through a total of three apertures before entering the pinch focus region. As a consequence only some 10-15 mA of the original 380 mA potassium beam is admitted to the adiabatic focus region.*

---

[2] A. Tauschwitz, M. de Magistris, E. Boggasch, M. Dornik, D.H.H. Hoffmann, J. Jacoby, W. Seelig, P. Spiller, H. Wetzler, "Performance of a Shape Optimized Conical Plasma Lens", GSI Report 94-10 (1994) 10.

Experimental results indicate that pinch focusing increases the peak current density of the beam by a factor of more than 70. Without focusing the beam size is determined by the inner diameter of the tapered glass tube and little current is transported to the Faraday cup as shown in Figure 1-8. With focusing the beam size is comparable to or smaller than the one millimeter moveable pinhole in front of the Faraday cup.

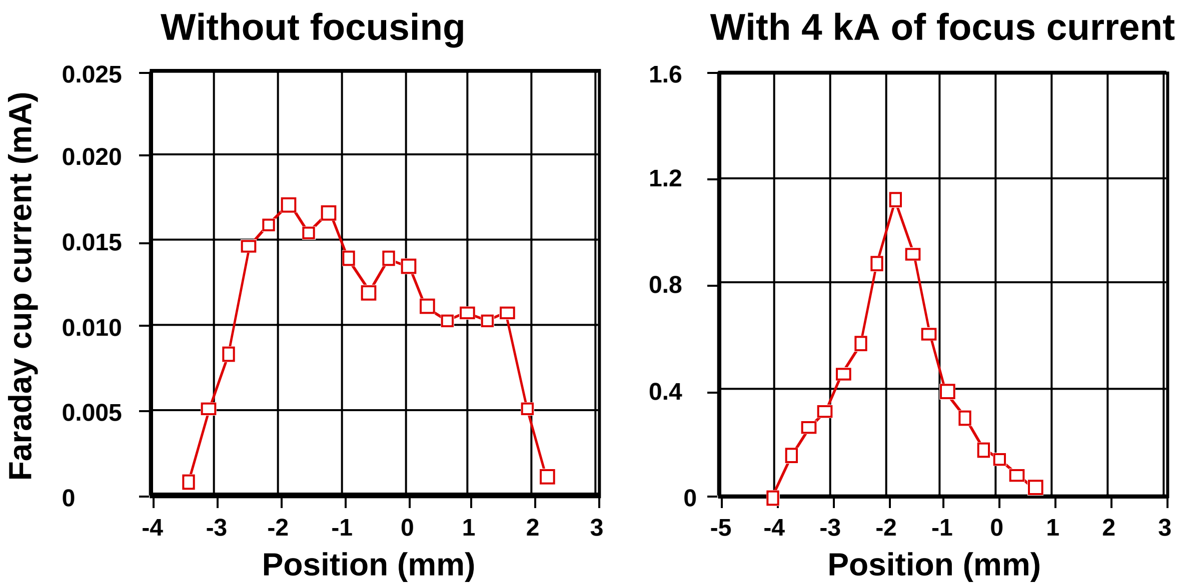

Figure 1-8*. Current detected by the Faraday cup as the 1 mm diameter pinhole is scanned across the beam without and with adiabatic focusing.*

Channels for Guiding Intense Ion Beams

In a target chamber the proper production and control of the Z-discharge channels is an essential element of establishing the feasibility of this concept. Even if targets illuminated by a single ion beam prove practical, using two current-carrying channels would probably be more efficient. As a practical minimum we are investigating techniques that bring two ion beams to the target, with two to four channels to return the current to the chamber wall. There are many issues that must have acceptable solutions. These include breakdown control within the target chamber; matching and synchronizing the discharge size and current to the incoming ion beam; understanding and avoiding potential channel and/or beam instabilities; etc. Many elements needed for generating proper Z-discharge channels must be integrated with elements of existing inertial fusion chamber concepts.

We are investigating a method of generating Z-discharge channels at parameters near those required in a power plant. A krypton-fluoride (KrF) laser is used to form a straight tenuous plasma approximately 5 mm in diameter and 440 mm long in a gas composed of 5 to 10 Torr nitrogen and 0.2 to 0.5 Torr benzene. The laser ionizes the benzene by a two-step photon absorption process. Energy from a prepulse capacitor is switched across the discharge electrodes at the moment the laser is fired. This energy flows through the plasma produced by the laser and strongly heats the gas. Expansion of the gas creates a low density bubble along the path of the laser that connects the discharge electrodes. The gas density in this bubble is as much as a factor of ten lower than ambient. After a delay of 10-30 μs the main capacitor bank is fired across the discharge electrodes forming a new discharge through the gas bubble. A voltage of 5 to 30 kilovolts from a capacitor bank is applied, causing a current of as much as 50 kA to flow along the channel. The experimental chamber is shown in Figure 1-9. The laser enters the chamber from the left and exits to the right through small holes in two copper plates that form the electrodes for the Z-discharge.

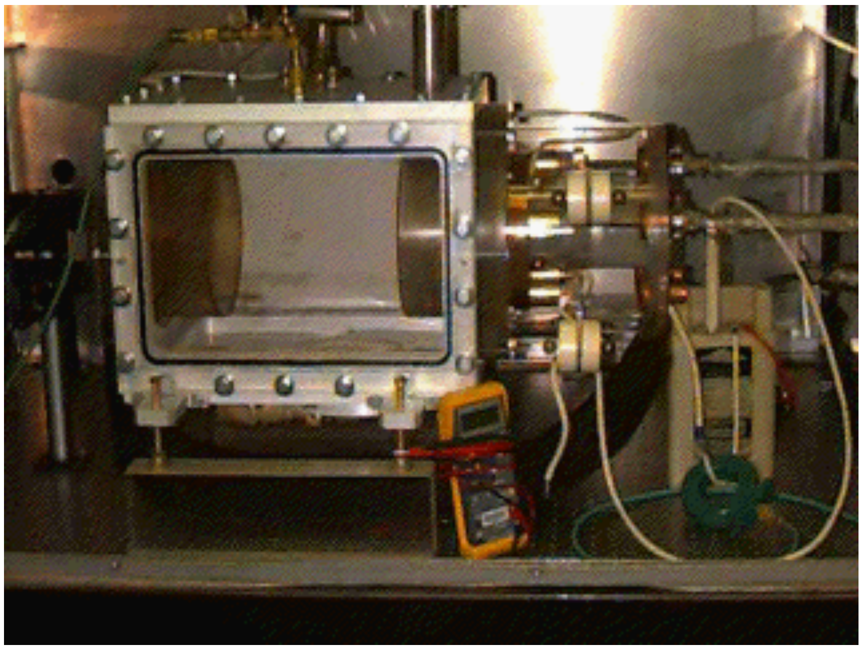

Figure 1-9. *The channel discharge apparatus; channels are formed over a maximum distance of 44 cm.*

This discharge can be precisely timed and could in principle easily be synchronized to an incoming heavy ion beam. Oscillograms showing the discharge timing are shown in Figure 1-10.

We are also investigating the interaction of Z-discharge currents in the vicinity of a fusion target using the model targets shown in Figure 1-11. A model target is mounted at the center of the discharge chamber and the KrF laser beam passes through a hole in the target. The KrF laser designates a channel from each discharge electrode to the model target. The needles visible in photo are used to control the current in the vicinity of the target.

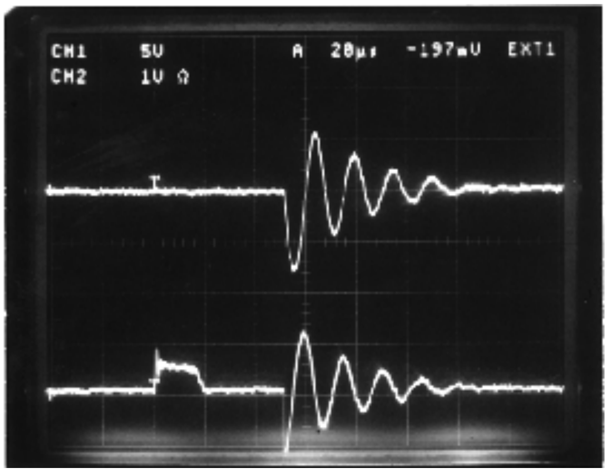

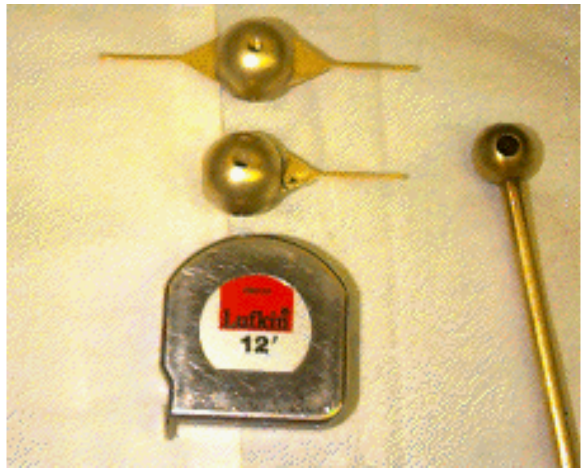

***Figure 1-10* (above left)**. *Oscillograms of the Z-discharge current and voltage. The upper trace shows the discharge current at 18.3 kA/div. The lower trace is the voltage across the discharge at 10 kV/div. The laser is fired at time T.* ***Figure 1-11.* (above right)** *Various models used to study and control the interaction of the discharge currents near a fusion target.*

A pseudo-color, 20 ns gated TV image of current-carrying channels intersecting a brass target is shown in Figure 1-12. Current flows vertically to the chamber wall through the two brass rods that support the ball. The image was obtained near the time the channel diameter was at a minimum. For this particular experiment, the model target was a 2 cm diameter ball without needles.

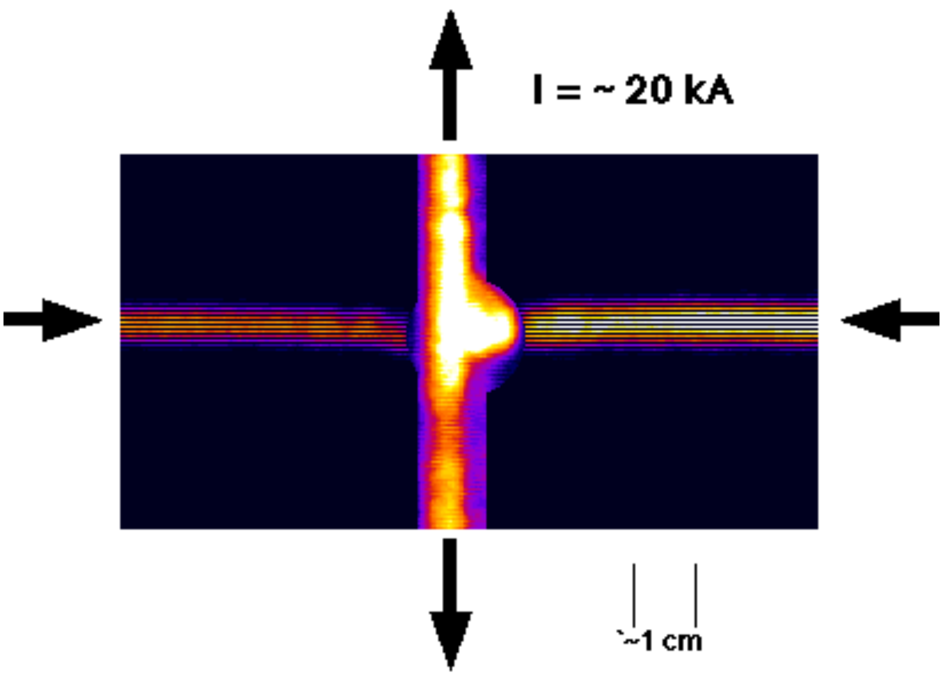

Figure 1-12*. Gated TV image of discharge channels intersecting a brass ball. The current is carried to the ball from the right and left.*

These experiments show that plasma channels 10 to 20 cm long and carrying kiloampere currents can be made to intersect at a model target and return to the wall as would be required in a heavy ion fusion power plant. These channels are approximately the correct diameter to match the beams to a fusion target. Moreover, the channels can be easily be synchronized to the incoming beams of heavy ions. More experiments, particularly with channels measured in meters, are required to establish this technique as a viable approach to channel transport in an operating inertial fusion power plant.

Reported by Thomas J. Fessenden

The Solenoidal Transport Option For IFE Drivers and the Near Term Facility

Solenoidal magnets have been used for the beam transport systems in all the high current electron induction accelerators that have been built in the past several decades. They have also been considered for the front end of the transport system for heavy ion accelerators for Inertial Fusion Energy (IFE) drivers, but this option has received little attention in recent years. The analysis reported here was stimulated mainly by the recent effort to define an affordable "Integrated Research Experiment" (IRE) that can meet the near-term needs of the IFE program. The 1996 FESAC IFE review panel agreed that an integrated experiment is needed to fully resolve IFE heavy ion driver science and technology issues; specifically, "the basic beam dynamics issues in the accelerator, the final focusing and transport issues in a reactor-relevant beam parameter regime, and the target heating phenomenology". The development of concepts that can meet these technical objectives and still stay within the severe cost constraints all new fusion proposals will encounter is a formidable challenge.

Solenoidal transport has a favorable scaling at high charge-to-mass ratio, low kinetic energy, and large beam radius (the main reasons why it is preferred for electrons in the region below 50 MeV). These facts were recognized in a recent LLNL conceptual study of high intensity induction linac-based proton accelerators for Accelerator Driven Transmutation Technologies, where solenoidal transport was chosen for the front end. Reducing the ion mass is an obvious scaling to exploit in an IRE design, and the output beam voltage will necessarily be much lower than that of a full scale driver, so solenoids will be considered as one option for this experiment as well. A testbed using solenoidal transport could in principle provide a high degree of flexibility, enabling experimental studies with a range of ion masses, since it has no required periodicity tied to the ion mass. This flexibility could be helpful in exploring a broad parameter range in the physics of target chamber transport and target heating, for example. It is also worth noting that a significant industrial capability exists to produce superconducting solenoidal magnets for Magnetic Resonance Imaging (MRI) machines; this capability could be exploited in our effort to find a cost-effective approach to an IRE.

The use of solenoids for an IRE would have limited appeal if the physics and technology were unrelated to a full scale driver. There is no question that most of a heavy ion driver's transport system (the high energy section) would consist of superconducting magnetic quadrupoles. Nevertheless a "front end" design based on solenoidal transport can form the basis of a significantly different overall driver architecture. With a solenoidal front end, we can eliminate beam combining and consider drivers with a smaller number of beam channels.

Constraints on final focusing are one reason solenoids have received little attention in recent years. Unneutralized ballistic focusing has been considered a low-risk focusing option for heavy ion fusion. This focusing option requires division of the driver output into multiple beams, of the order of 32, and particle energies of 10 GeV or more. In contrast solenoids are more attractive for systems that have a few large beams and relatively low kinetic energy. Because unneutralized focusing restricts us to high kinetic energy, we have, for several years, been studying and experimenting with beam focusing and transport methods that employ neutralization. The goal of this research is to allow us to use ions with higher charge-to-mass ratios, lower kinetic energy, and perhaps fewer beams. These steps may lead to less expensive driver systems. Two experiments on neutralization are described in this report. An implicit assumption in considering accelerator architectures with low kinetic energy and a small number of beams is that at least one of the neutralized focusing and transport schemes will be viable. Alternatively, if beam splitting prior to final focus proves feasible, acceleration of a small number of beams can still be used even with a “conventional” unneutralized final focus system.

Beam quality (transverse and longitudinal emittance) is another consideration that has favored high mass, high kinetic energy and large numbers of beams. In this regard also, other focusing methods may be more favorable than the unneutralized ballistic method. These alternative focusing systems may be particularly useful for systems using solenoids because solenoidal channels transporting a high line charge density during acceleration have relatively higher risk of emittance dilution from space charge effects.

As mentioned, the essential feature of solenoidal transport which makes it an attractive option is its high current per beam at low energy. The line charge density of a space-charge-dominated nonrelativistic beam of radius (a) transported in a continuous solenoidal is given by

$$\lambda = \frac{\pi\varepsilon_0 qeB^2a^2}{2M} \qquad (1)$$

where B is the magnetic field, M is the ion mass, and qe is the ion charge. The model is that of a uniform density beam injected into the solenoidal field from a field-free ion source, the well known Brillouin flow condition that corresponds to the maximum line charge density that can be transported for that magnetic field strength. A similar relation holds for a series of solenoidal lenses that have a periodic axial variation in B on a scale length small compared to the cyclotron wavelength, if we replace B^2 by its average value. This situation applies, for example, to interruptions in the solenoids from the accelerating gaps. Numerically,

$$\lambda = 10\left(\frac{133}{A}\frac{q}{}\right)B^2a^2\ \mu\mathrm{C/m} \qquad (2)$$

where A is the ion mass in amu, B is in tesla, and “a” is in meters. From this formula, it appears that heavy ion beams with $\lambda \approx 10\ \mu\mathrm{C/m}$ are transportable using high fields and large beam radii.

The radial potential drop across the beam is $\delta\phi \equiv \lambda/4\pi\varepsilon_0$, and the ratio of $\delta\phi$ to the beam edge voltage V is equal to the dimensionless perveance, Q. For $\lambda \approx 10\ \mu C/m$ this yields $\delta\phi \equiv 90 kV$ In Brillouin flow, the axial velocity is a constant across the beam, equal to

$$v_Z^2 = v_0^2 - qe\lambda/2\pi\varepsilon_0 M \qquad (3)$$

, where $Mv_0^2/2 = qeV$. Since the beam current is given by $I = \lambda v_Z$, it is easy to show that there is a maximum value of the current that can be transported, corresponding to a maximum perveance of Q=2/3. This limiting current can be a significant factor in determining the minimum gun voltage for injection into a solenoidal channel designed to carry very high line charge densities.

By comparison, in a continuous electrostatic quadrupole channel, the maximum line charge that can be transported is independent of the ion mass and the beam voltage. It depends on the voltage between adjacent electrodes φ_{es}, the longitudinal occupancy fraction of the quadrupole field η, the radius of the beam a, and the radius of the pole-tip b.

$$\lambda \approx \frac{\pi\varepsilon_o\eta\varphi_{es}}{2}\frac{a^2}{b^2} \qquad (4)$$

For a typical optimized design, the bore radius is approximately 2 cm and λ is approximately 0.25 µC/m. Due to the small beam size the current density averaged over the area of an entire array is comparable for magnetic solenoids and electrostatic quadrupoles.

This electrostatic-quadrupole scaling leads to systems that have relatively large numbers of small beams. Thus the two main options for the front end of a driver are (1) a large number of small beams focused by quadrupole arrays or (2) the solenoid option described in this paper which employs a smaller number of larger beams.

The use of solenoids for high current ion transport raises several areas of concern which are now being studied in the Fusion Program. One of these is the effect of fringe field aberrations at the acceleration gaps, which are predicted to cause some degree of emittance growth. This is being studied with an analytical cold fluid model of the beam and with a new simulation code SALT, which treats non-linearity of beam and solenoidal fields. A second concern is the design of an ion source which could produce high enough currents to make efficient use of solenoids. We have made substantial progress in this area as described elsewhere in this chapter. Stability limits for transportable current are understood in a macroscopic sense at present. However the possibility of microinstabilities which limit emittance or current will be examined using SALT and our 2-d and 3-d particle codes (e.g., WARP, originally developed at LLNL).

Reported by Edward Lee (LBNL) and Richard Briggs (SAIC)